Intel Research

Proactive Computing

www.intel.com/research

Proactive computers will anticipate our needs and take action on our behalf.

Driving Toward a Future of Proactive Computing

A new era of computing is on the horizon. In this new era, billions of computing devices will be deeply embedded within our physical environment. These tiny sensors and actuators will silently serve us, acquiring and acting on a multitude of data to improve our lives and make us more productive. However, today's model of interactive computing, whereby individuals interact one-on-one with computers, inputting commands and waiting for responses, will not scale to the new environment.

The interactive computing model is already showing its limitations, as we begin to confront the challenges of dealing with multiple computers, from desktop and laptop systems to cell phones, PDAs, and a growing variety of consumer electronic devices. When we each have hundreds or thousands of devices to deal with, it will be impossible for us to interact directly with each one.

The time has come to transition from interactive to *proactive* computing. Proactive computers will anticipate our needs and take action on our behalf. Instead of serving as glorified input/output devices, humans will be freed to focus on higher-level functions.

The challenge for the research community is to explore the implications and issues associated with having hundreds or thousands of networked computers per person, and to overcome the significant obstacles involved in realizing the vision of proactive computing.

Challenges on the Road to Proactive Computing

We've identified three steps that are essential to making proactive computing a reality. The first is *getting physical* — connecting billions of computing devices directly to the physical world around them so that human beings are no longer their principal I/O devices. The next step is *getting real* — having computers running in real time or even ahead of real time, anticipating human needs rather than simply responding to them; early examples include airbags and antilock brakes. The third step is *getting out* — extending the role of computers from the office and home into the world around us and into new application domains.

In moving successfully through these critical steps, researchers will face difficult challenges. We have identified several key challenges and have initiated research projects to address them, drawing on the expertise of some of the very best people, both within our own laboratories and through our extensive network of university connections worldwide.

Research Challenges	Research Projects
Make It Personal	Ethnography, Distraction Free, Security and Privacy
Closing the Loop	Complex Adaptive Systems, Agent Negotiation
Anticipation	Speculative Execution, Machine Learning, Autonomy
Dealing with Uncertainty	Statistics, Robotics, Machine Vision
Macroprocessing	IPP, XML, Migration, Novel Storage, Live Databases
Deep Networking	802.11, UPnP, UWB, Ad-Hoc Networks, Tiny OS
Getting Physical	MEMS, Biotech, Antennas, Energy

Getting Physical

Today, human beings are the primary source/synch of data. To move beyond that role, computers must be better connected to the world around them and to each other. With that in mind, researchers are focused on new types of sensors and actuators, especially Microelectromechanical Systems (MEMS) and biochips, using our precision capabilities to drive MEMS to their limits. The goal is to enable computers to sense things that are dry and solid state and things that are wet, such as biological materials. This will vastly expand the range of places where we can install and network computers, and it will enable a variety of important wetware applications in areas such as health, pharmaceuticals, chemicals, and refineries.

At a more fundamental level, getting physical with billions of devices will relieve human beings of burdensome I/O devices and vastly expand the flow of raw data into information systems.

"MEMS technology enables the microscopic sensors and actuators that will play an essential role in connecting our physical world with our digital world."

Researcher:
Valluri (Bob) Rao, MEMS

Intel Fellow, Technology and Manufacturing Group

Group Director, Analytical and Microsystems Technologies

"We can transform health care through early disease detection and new forms of health care delivery. Our project uses Intel's advanced technology to enable a first generation of diagnostic chips."

Researcher:
Andrew Berlin, Precision Biology

Principal Engineer, CQN Microsystems and MEMS, Technology and Manufacturing Group

Small, inexpensive, low-powered sensors and actuators will enable us to achieve a hundredfold increase in the size of the Internet beyond the growth we're already expecting.

Deep Networking

Small, inexpensive, low-powered sensors and actuators, deeply embedded in our physical environment, can be deployed in large numbers, interacting and forming networks to communicate, adapt, and coordinate higher-level tasks. As we network these micro devices, we'll be pushing the Internet not just into different locations but deep into the embedded platforms within each location. This will enable us to achieve a hundredfold increase in the size of the Internet beyond the growth we're already anticipating. And it will require new and different methods of networking devices to one another and to the Internet.

Wireless technology is one requirement for deep networking. Intel is already addressing this requirement in its 802.11 wireless products and is aggressively driving research into affordable and ad hoc wireless networking, especially the networking of sensors. We're exploring ultra wideband radio, a new networking technology that will provide greater density and bandwidth, and enable more ways of going about our business. We're also enabling new types of antennas, and new ways of integrating radios with MEMS and CMOS devices.

"As the billions of tiny embedded devices become networked, they will form the vast majority of the Internet. They present radically different networking requirements and offer tremendous room for innovation."

Researcher:
David Culler,
Extreme Networked Systems

Director, Intel Research Berkeley

Professor, Computer Science,
University of California Berkeley

"The synthesis of RF and digital computation will be a necessary part of realizing the proactive computing vision. UWB and related technologies offer opportunities to bring Moore's Law to the radio world to deliver high bandwidth, low power, highly integrated devices."

Researcher:
Kevin Kahn, Ultra Wideband Radio

Intel Fellow, Corporate Technology Group

Director, Communications and
Interconnect Technology

"The performance, capacity, and economies of scale of today's CMOS process technology allow, for the first time, the infusion of low cost wireless technologies that can span a range of applications from remote sensing devices to powerful base station and access points."

Researcher:
Stephen Pawlowski, CMOS Radio

Intel Fellow, Corporate Technology Group

Director, Communications and
Interconnect Technology

With thousands of devices, many of which will be resource impoverished, it will be necessary to provide a fluid mechanism for computations to move from one node to another.

Macroprocessing

Once computing devices are both connected and networked, we'll need to enable macroprocessing. Software must work across a wide range of diverse platforms and networks. These challenges require a variety of research efforts.

- **Intel® Integrated Performance Primitives (IPP) –** This technology will ensure that core kernels of software run efficiently, no matter which Intel architecture the software is running on. The core innermost loops can easily be optimized and ported from architecture to architecture.
- **XML-Based Capabilities –** In effect, XML provides a universal type library and labeling mechanism that frees data and programs to move across organizations, across software platforms, across many different types of enterprise architectures and software.
- **Migration –** With thousands of devices, many of which will be resource impoverished, it will be necessary to provide a fluid mechanism for computations to move from one node to another. This is a challenge that needs to be solved in a highly heterogeneous environment, both in terms of platforms and connectivity. It needs to support people as they move from place to place as well as adapting to the power and bandwidth limitations of the devices they use. In Pittsburgh, we are working on a mechanism called Internet Suspend/Resume that enables execution to be suspended at one Internet site and seamlessly resumed at another. We are working with researchers at the University of Washington developing the virtual machine and middleware support needed for robust migration.
- **Storage –** The new edge nodes will have new types of storage and file systems. It will be a challenge to connect those devices with our personal servers, and then to connect the servers, using peer-to-peer computing, to form totally global, persistent, secure file stores.
- **Querying "Live" Data –** Once programming and storage in the large are enabled, database-like queries, using live data from sensors embedded throughout the environment, will be possible. Research into live query tools is underway at the University of California-Berkeley in a project called Telegraph. Similar work is underway at Cornell and elsewhere.

With the launch of their recent efforts on autonomic computing, our research colleagues at IBM are undertaking an enormously important piece of the proactive computing challenge. Making systems self-configuring, self-managing, and self-healing will provide a firm foundation on which to build. Given a macroprocessing layer at which systems can organize and manage themselves, we can move forward to the higher level problems associated with proactive computing — dealing with uncertainty, anticipating user requirements, acting on our behalf, and empowering individuals to attain their goals.

"The Ubiquity project focuses on the Personal Server concept, a device that overcomes many of the limitations of today's mobile computers, such as their small screens, large size and cumbersome weight. Carry and access all your digital data, all the time, wirelessly through nearby displays and keyboards."

Researcher:
Roy Want, Ubiquity

Principal Engineer,
Intel Research,
Corporate Technology Group

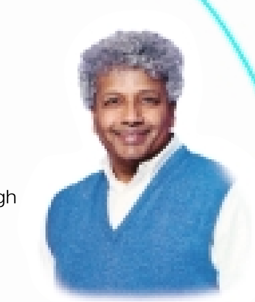

"Proactive techniques are important for providing failure-resistant, scalable, and secure access to shared information by mobile and static users over the wireless and wired networks of the future."

Researcher:
Mahadev Satyanarayanan,
Ubiquitous Storage

Director, Intel Research Pittsburgh

Carnegie Group Professor,
Computer Science,
Carnegie Mellon University

Excess computation and communication capacity will be harnessed to fetch and manipulate information, producing answers before they are required.

It's crucial that we take that empowerment in personal computing with us as we drive forward into this proactive era.

Dealing with Uncertainty

Increased access to data is a necessary, though not sufficient, condition if we are to wring ever-greater productivity gains from human endeavors. As we move up the stack and integrate data from the world around us, we must learn to deal with uncertainty. The physical world does not exhibit the deterministic behaviors computer scientists have come to know, model, and love.

As a community, we have to learn how to use statistical methods much more effectively. This is already happening to some extent, with Google* being the most widely used example. The researchers who are achieving best-of-class performance in areas like computer vision, speech, and data mining are using Hidden Markov Models, Bayesian Models, and other statistical techniques. We are investing in broadening this competency throughout our organization and we encourage other organizations to do the same.

"One of the challenges for computer vision is to handle the complexity of the real world. We need to close the gap between simplistic models and the physical reality."

Researcher:
Horst Haussecker,
Computational Nanovision

Senior Researcher,
Microprocessor Research, Intel

"Statistical methods exploit the fact that measurements of real-world things have varying values, not just a single unique value. In the past, the game was to find the mean and stick to it. Now, the game is to flow with the variance."

Researcher:
Gary Bradski, Statistical Models in Manufacturing and Vision

Senior Researcher,
Microprocessor Research, Intel

Anticipation

In the era of proactive computing, software will anticipate our needs. Excess computation and communication capacity will be harnessed to fetch and manipulate information, producing answers before they are required — much as a chess champion predicts his opponent's moves many steps into the future. We see glimpses of anticipation today, both in the speculative execution features of Intel's processors and in some of the most advanced web proxy engines.

A development we find particularly exciting is the emergence of a new genre of machine learning tools that is firmly grounded in statistical methods. Systems such as those under development by Koller (Stanford), Fox (University of Washington) and Thrun (Carnegie Mellon University), exploit uncertainty to support robotic hypothesis generation, a key stepping stone to anticipation.

"We are inventing wellware — that is, the software and hardware to help people stay healthy — by conducting fieldwork into the daily lives of elders with Alzheimer's, physical disabilities, cancer, and cardiovascular disease. Our ultimate goal is to help people 'age in place' from their own homes through systems that work contextually, reliably, and proactively on their behalf."

Researcher:
Eric Dishman,
Proactive Consumer Health

Principal Investigator,
Intel Research,
Corporate Technology Group

"We've coupled system dynamics with visualization, using modeling and simulation to explore the dynamics, interactions, and relationships in business, and then added tangible interfaces to create new interactive real-time mechanisms."

Researcher:
Mary Murphy-Hoye, Supply Chain Visualization

Director, IT Industry Research

IT Principal Engineer

Closing the Loop

Once we've learned to manage uncertainty and can get computers to anticipate our needs, the next challenge is bridging the gap between anticipating those needs and acting on them — albeit under human supervision and in a predictable way. We refer to this challenge as closing the loop.

The difficult step here is not so much in determining the action to be taken, but to develop the feedback and control mechanisms essential to the stable operation of any closed loop system. Today, most systems rely on human beings to close the loop and provide stability, placing human beings under intolerable stress.

To better understand the challenge, we've been working with our colleagues at the MIT Sloan School on modeling Intel's supply chain — a series of nested control loops that must be monitored and managed carefully. At the Internet scale, we must address the use of software agents to close loops. The challenge, of course, is not to close a single loop but to engineer systems in which billions of computers, and an unlimited number of software agents, are simultaneously closing many intertwined loops.

How do we get all these computers acting in a predictable manner on our behalf? How many agents can we have working for us at a given time, without demanding so much of our attention? We've been working with our colleagues at the Santa Fe Institute on these complex networks, and we're beginning to see significant progress in this area.

Make It Personal

As we move to the top of the stack, we focus on human empowerment, on making computing personal. A strong sense of individual empowerment has helped drive the growth of our industry, and it's crucial that we take that empowerment in personal computing with us as we drive forward into this proactive era. We also must manage security and privacy concerns — not just for a single, homogeneous society, but for a world filled with different cultures with differing expectations for security and privacy.

To begin to address this challenge, our Intel labs employ a world-leading group of ethnographers and social scientists who spend their time studying people in small numbers, analyzing in detail what they're doing with technology, why they're doing it, how it can help them in their lives, and how it can improve their quality of life.

"All these objects will be communicating; therefore, we need to develop ways to have them autonomously configure and maintain themselves. People have more important things to do."

Researcher:
Gaetano Borriello,
Distraction-free Systems

Director, Intel Research Seattle

Professor,
Computer Science & Engineering,
University of Washington, Seattle

"In order to deliver on the promise of proactive computing, we need to truly understand not only what users might want in the future, but what real people are doing now. And not just in the familiar expected places, but all over the world."

Researcher:
Genevieve Bell, Ethnography

Principal Investigator,
Intel Research, Corporate Technology Group

Making the Vision a Reality

The entire proactive computing agenda has to focus on continuing the revolution in human empowerment that we started collectively 20 years ago. Intel is strongly committed to maintaining its focus on personal computing as we learn to harness hundreds or thousands of computers per person. We hope that you will join us in this adventure.

Printed in the USA/0502/ID/JH/5K

251214-001